

Noise Suppression - Decision Sheet

Classical vs neural, the four named approaches, what to pick, and the voice-vs-noise trade-off.

Two families

Classical (subtract)	Measure noise in the gaps, subtract it	Steady noise only; leaves musical noise
Neural (keep voice)	Pre-trained model keeps learned voice	Removes sudden noise; costs CPU/GPU + model

The four named approaches

WebRTC built-in NS	Classical, free, in every browser	Fan / AC; misses barks and clatter
RNNoise	Hybrid GRU, ~85 kB, any CPU	~60x real time x86; ~10 ms look-ahead; BSD
Krisp (SDK)	On-device neural, Small + Big models	No audio upload; ~25 ms latency; near-field
NVIDIA noise removal	GPU model (Maxine / Broadcast)	Bigger model off the CPU; needs NVIDIA GPU

Choose by use case


General meeting, mixed devices	WebRTC NS, or RNNoise
Contact centre (open-plan)	On-device neural CPU model
Telemedicine / finance (privacy)	On-device neural only
Podcast / recorded, no time limit	Largest neural model, offline
Streamer on a GPU workstation	NVIDIA noise removal
Embedded / low-power device	Classical, or RNNoise if it fits
Music / instrument capture	Turn suppression OFF

The voice-vs-noise trade-off

- SNR = how far the voice (dB) sits above the noise (dB). Calls sit nearer 20 dB than 0 dB.
- More suppression removes more noise AND more voice - it is never free.
- Example: voice at 10 dB SNR, pull noise down 18 dB -> effective SNR ~28 dB (comfortable).
- Push too hard and consonants (t, k, s) get cut: lower noise, muffled voice - a bad trade on a call.
- Judge it with ITU-T P.835: SIG (voice damage), BAK (residual noise), OVRL (overall).

Remember

- ☐ Classical suppression handles steady noise for free; it cannot remove a bark, a door, or other voices.
- ☐ Neural suppression has learned what voice looks like, so it removes sudden non-stationary noise.
- ☐ RNNoise is a small ~85 kB hybrid model that runs many times faster than real time on a plain CPU.
- ☐ Krisp runs on-device (no audio upload) with Small/Big models; NVIDIA runs a bigger model on the GPU.
- ☐ Match the model size to the device: a Big model stutters on a thin laptop or phone.
- ☐ Test with the P.835 SIG/BAK/OVRL split so a good noise number is not hiding voice damage.



GO DEEPER · VOLUME 4 OF 7

Audio for Video: A Complete Guide to Sound in Video Products

From PCM and Codecs to Loudness, Accessibility, and Real-Time Voice Processing

By Nikolay Sapunov, CEO at Fora Soft · 1,229 pages · Kindle & paperback on Amazon

[Read it on Amazon →](#)
amazon.com/dp/B0HH6PRZH1

This sheet compares classical and neural suppression, the four named approaches and the voice-versus-noise trade-off. Volume 4 covers real-time voice processing in depth: noise suppression, echo cancellation, AGC and VAD, plus the codecs and loudness rules the cleaned voice goes into.

WHAT YOU WILL FIND INSIDE

- Noise suppression in the real-time voice chain
- Echo cancellation, AGC and VAD
- Opus and speech codecs after suppression
- Loudness targets and LUFS for voice

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI

Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF

Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery

Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing

Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs

Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video

Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Choosing noise suppression for your calling app?

We will test WebRTC NS, RNNoise and a neural SDK on your real noise and devices and recommend the setup that keeps voices clear. 250+ video products shipped since 2005.

[Book a 30-min call →](#)
[WhatsApp →](#)
[Email us →](#)